To Keep or not to Keep: The Problem of Selecting Pictures from Personal Digital Collections

Andrea Ceroni

Data Scientist at SmartRecruiters GmbH
a.ceroni@smartrecruiters.com

18.05.2019
Ravenna, Italy

Leibniz Universität Hannover

ForgetIT

www.forgetit-project.eu

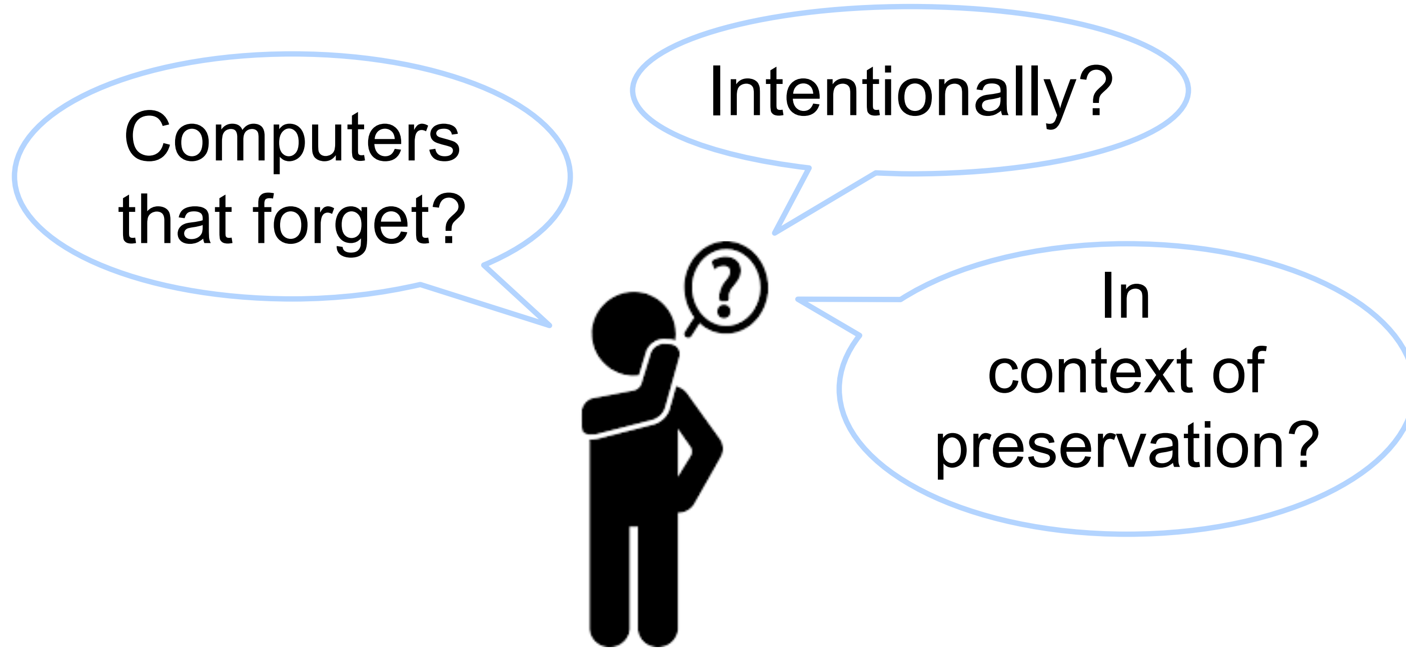

However, nowadays we are facing:

- dramatic increase in content creation (e.g. digital photography)
- increasing use of mobile devices with restricted capacity
- inadvertent forgetting (loss of data) due to lack of systematic preservation

Forgetting plays a crucial role for human remembering and life in general

- Forgetting of irrelevant details
- Focus on important information

Shouldn't there be something like forgetting in digital memories as well? →

Invest in preserving just the important information

Scenario

Personal Photo Explosion

- Photo taking is fun, effortless, and tolerated nearly everywhere
- Hundreds of pictures taken during vacations, trips, ceremonies…

What to best do with all of these photos?

How to select important photos for future revisiting and preservation?

Problems

High User Investment

- Great effort in revisiting, annotating, organising, making summaries
- Such effort increases with the size of the collections

Personal Collections become “Dark Archives”

- Photos are moved to some storage device
- Photos are rarely accessed and enjoyed again

Meeting user expectations

- What are the photos important to the user?
- What makes a photo important?
- Presence of personal (and hidden) attachment due to memories and context

User Study

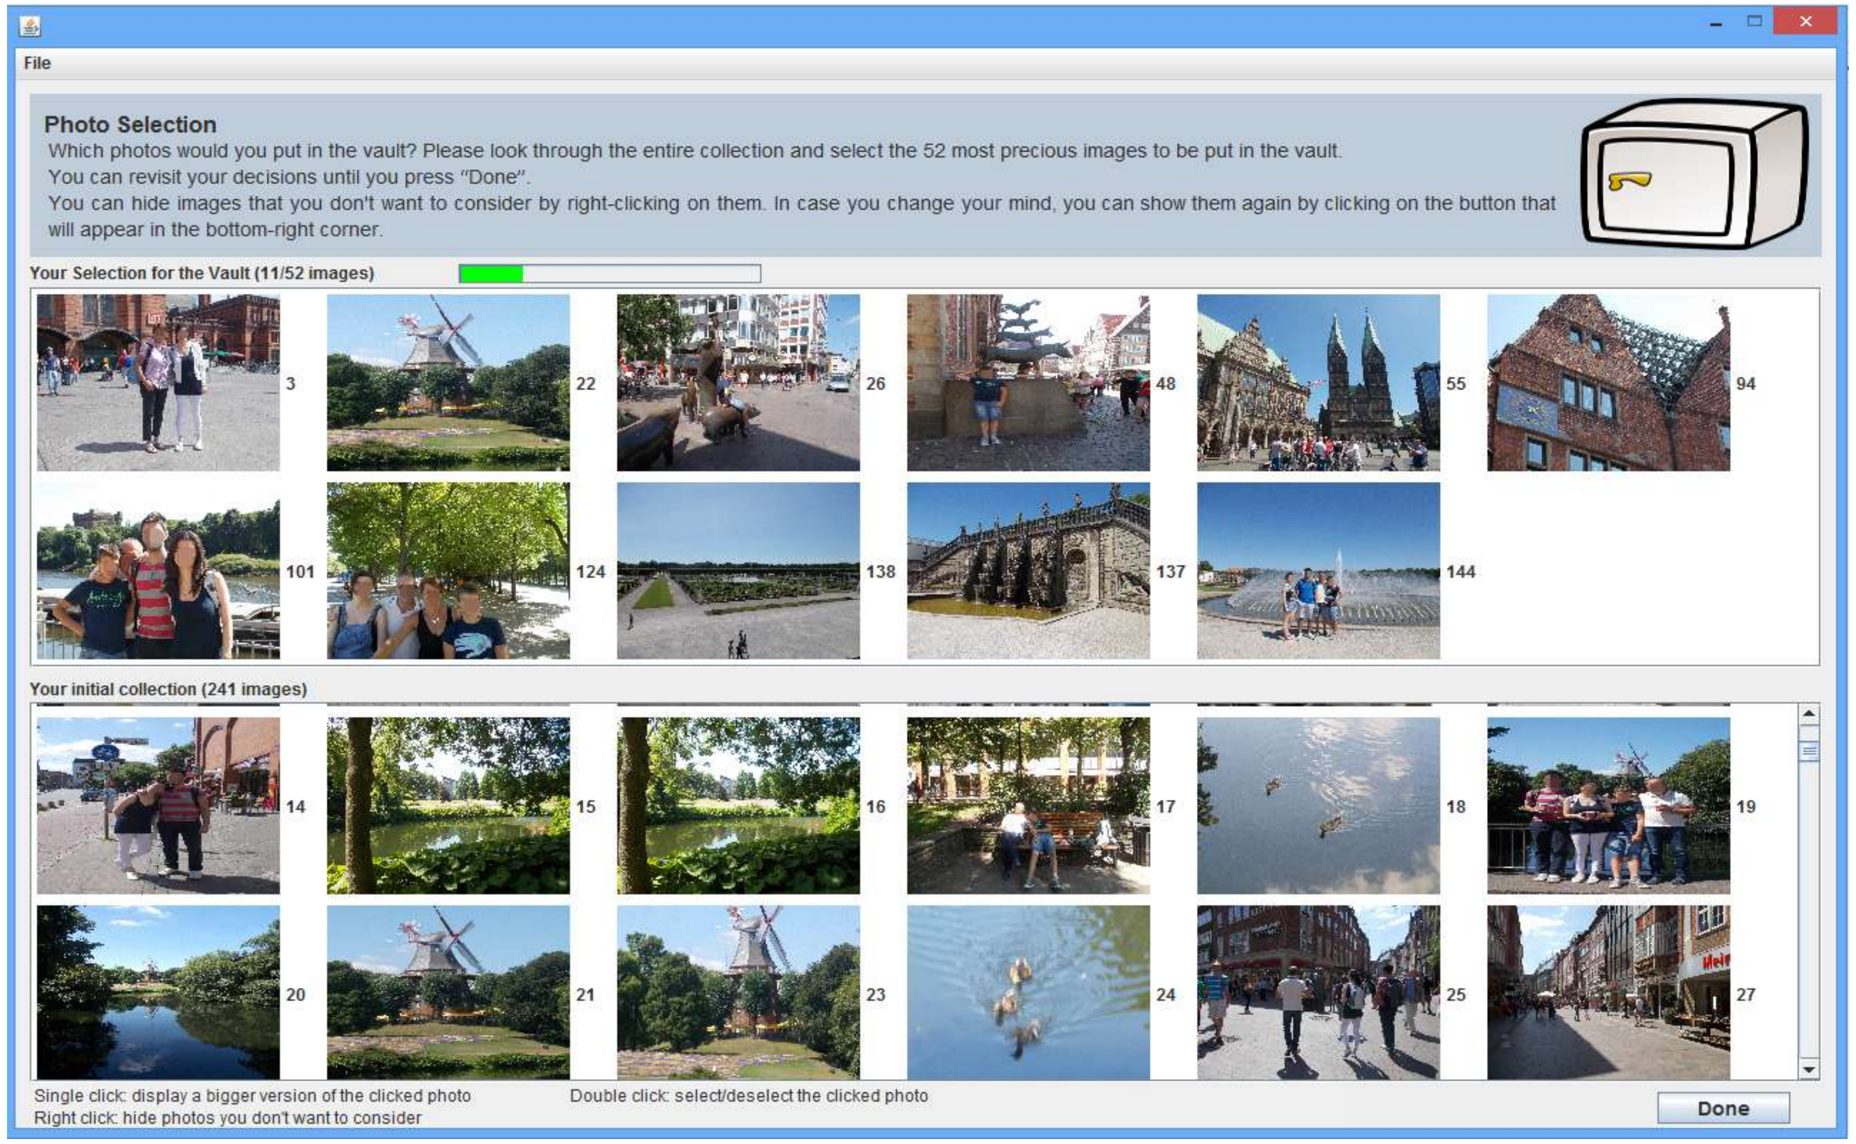

- ➢ Participants
  - ○ 42 people
  - ○ 91 collections
  - ○ 18,147 photos
- ➢ Task definition
  - ○ Each user provides one or more photo collections of personal events
  - ○ Selecting 20% of photos from each collection for preservation and revisiting
- ➢ Insights
  - ○ Image quality least important selection criterion
  - ○ Subjective and hidden aspects rated high
  - ○ Event coverage also highly important

Ceroni et al.. Investigating Human Behaviours in Selecting Personal Photos to Preserve Memories. In ICME Workshops 2015, pp. 1-6.

Automatic Photo Selection

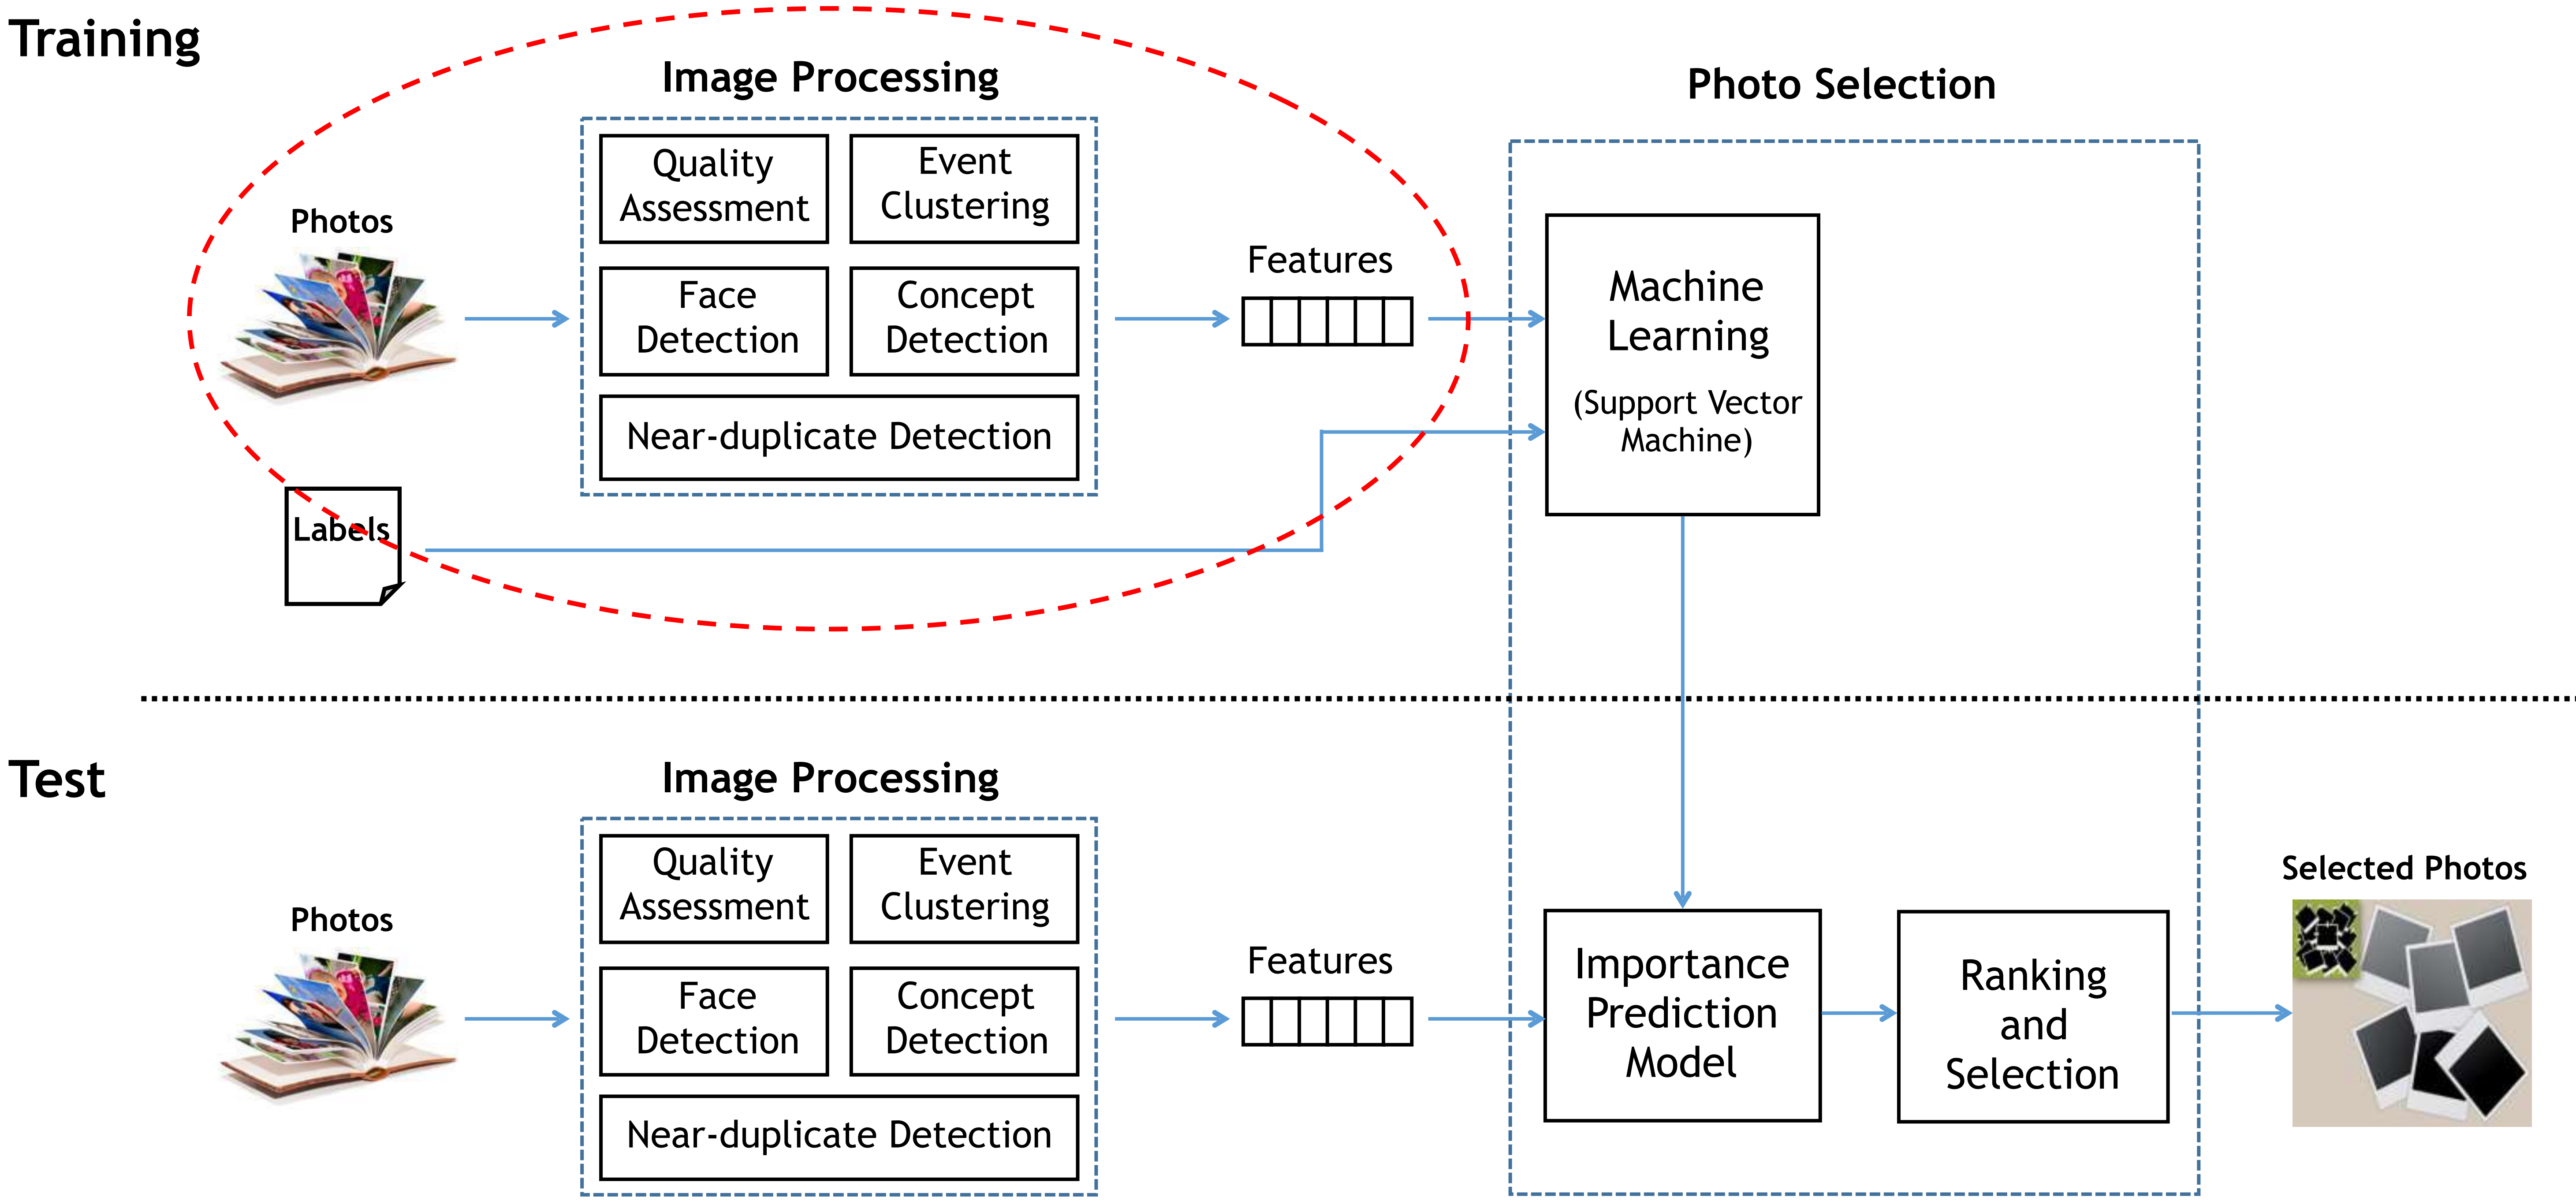

Ceroni et al. To Keep or not to Keep: An Expectation-oriented Photo Selection Method for Personal Photo Collections. In ICMR 2015, pp. 187-194.

Quality-based Features

Blur, contrast, darkness, noise

	Left photo	Right photo
Blur	0.533219	0.241118
Contrast	0.157777	0.107511
Darkness	0.870238	0.433792
Noise	0.179392	0.167515

Face-based Features

Presence, position, relative size of faces in each of 9 quadrants

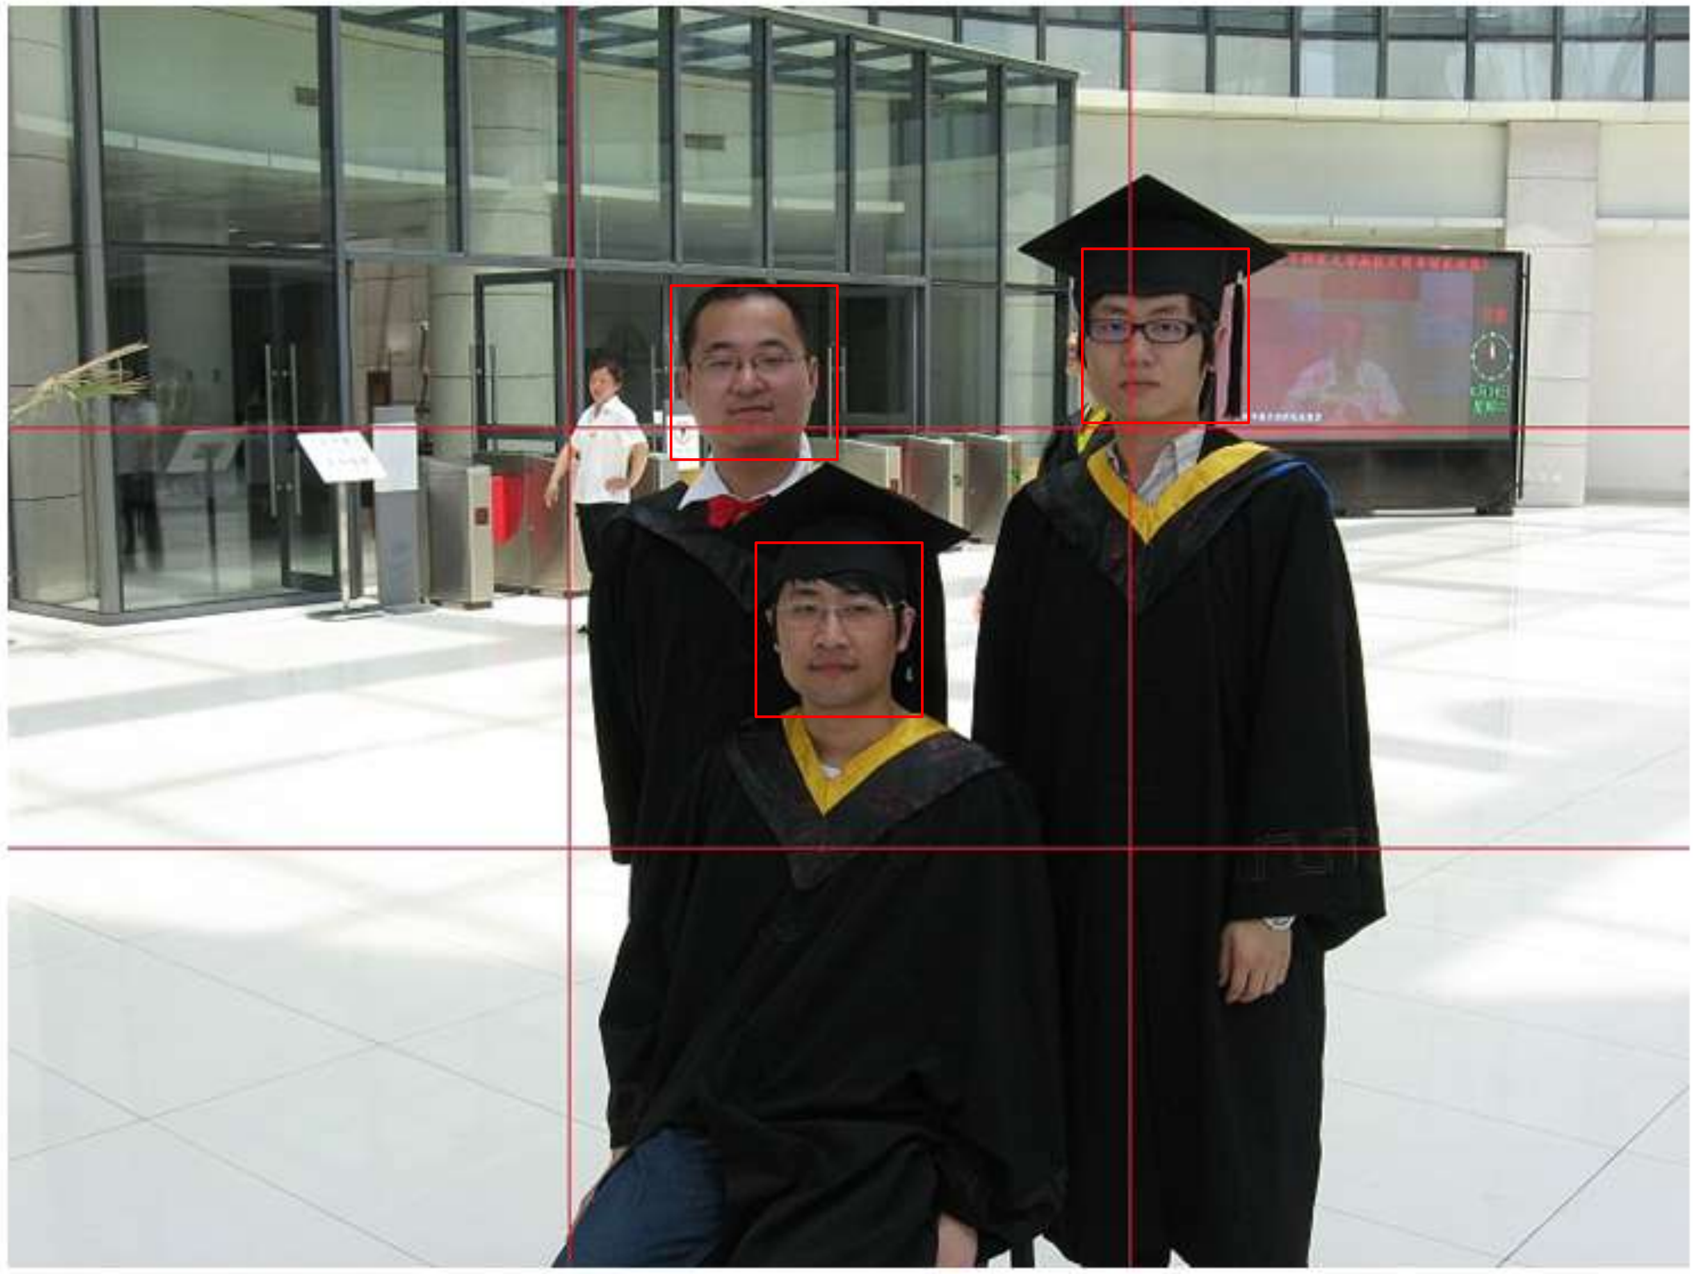

IMG_2440.JPG 3 N 0.009698 NE 0.009187 C 0.010998

Concept-based Features

346 concept detectors represented by SVMs (800 hours of video for training)

Top 10 concepts
- Outdoor: 0.9138
- Vegetation: 0.9
- Three_or_more_people: 0.89013
- Trees: 0.85785
- Building:0.83941
- Street: 0.81051
- Person: 0.79659
- Windows: 0.79222
- Sky: 0.76782
- Female: 0.75522

Collection-based Features

Temporal Clustering

group images belonging to the same sub-event

Near-duplicate Detection

identify similar shots of the same scene

Information about the clusters and near-duplicate sets each image belongs to:

- Size of its cluster
- Quality of its cluster
- Faces in its cluster
- Has near-duplicates?
- Size of its near-duplicates set

Automatic Photo Selection

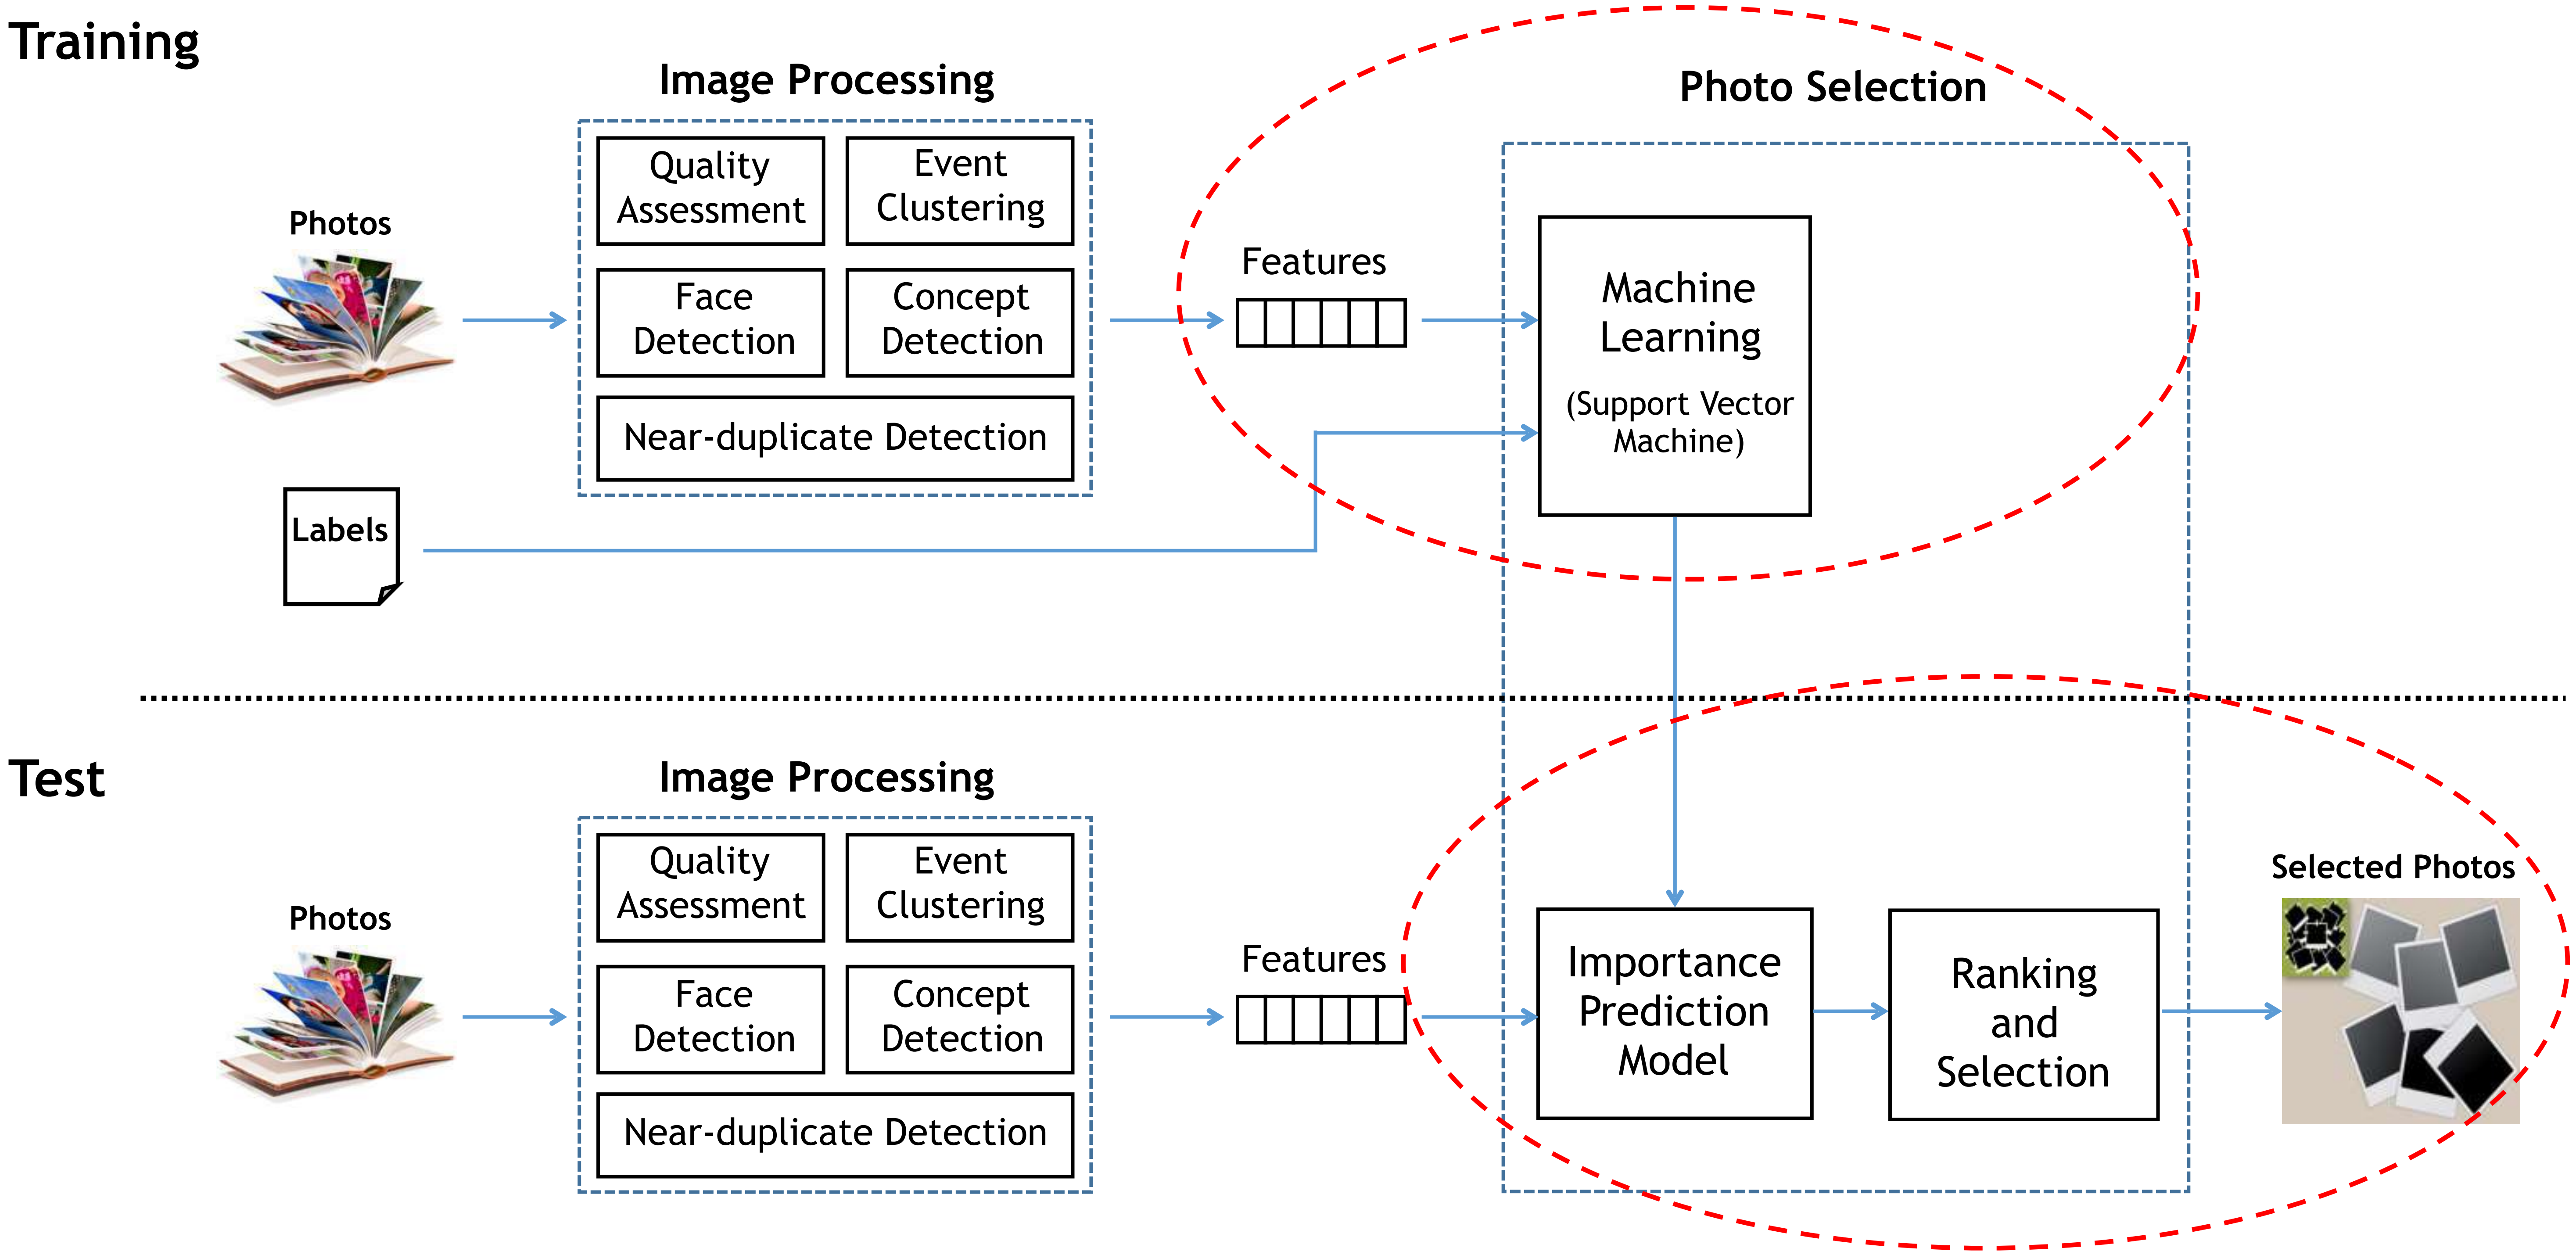

Importance Prediction

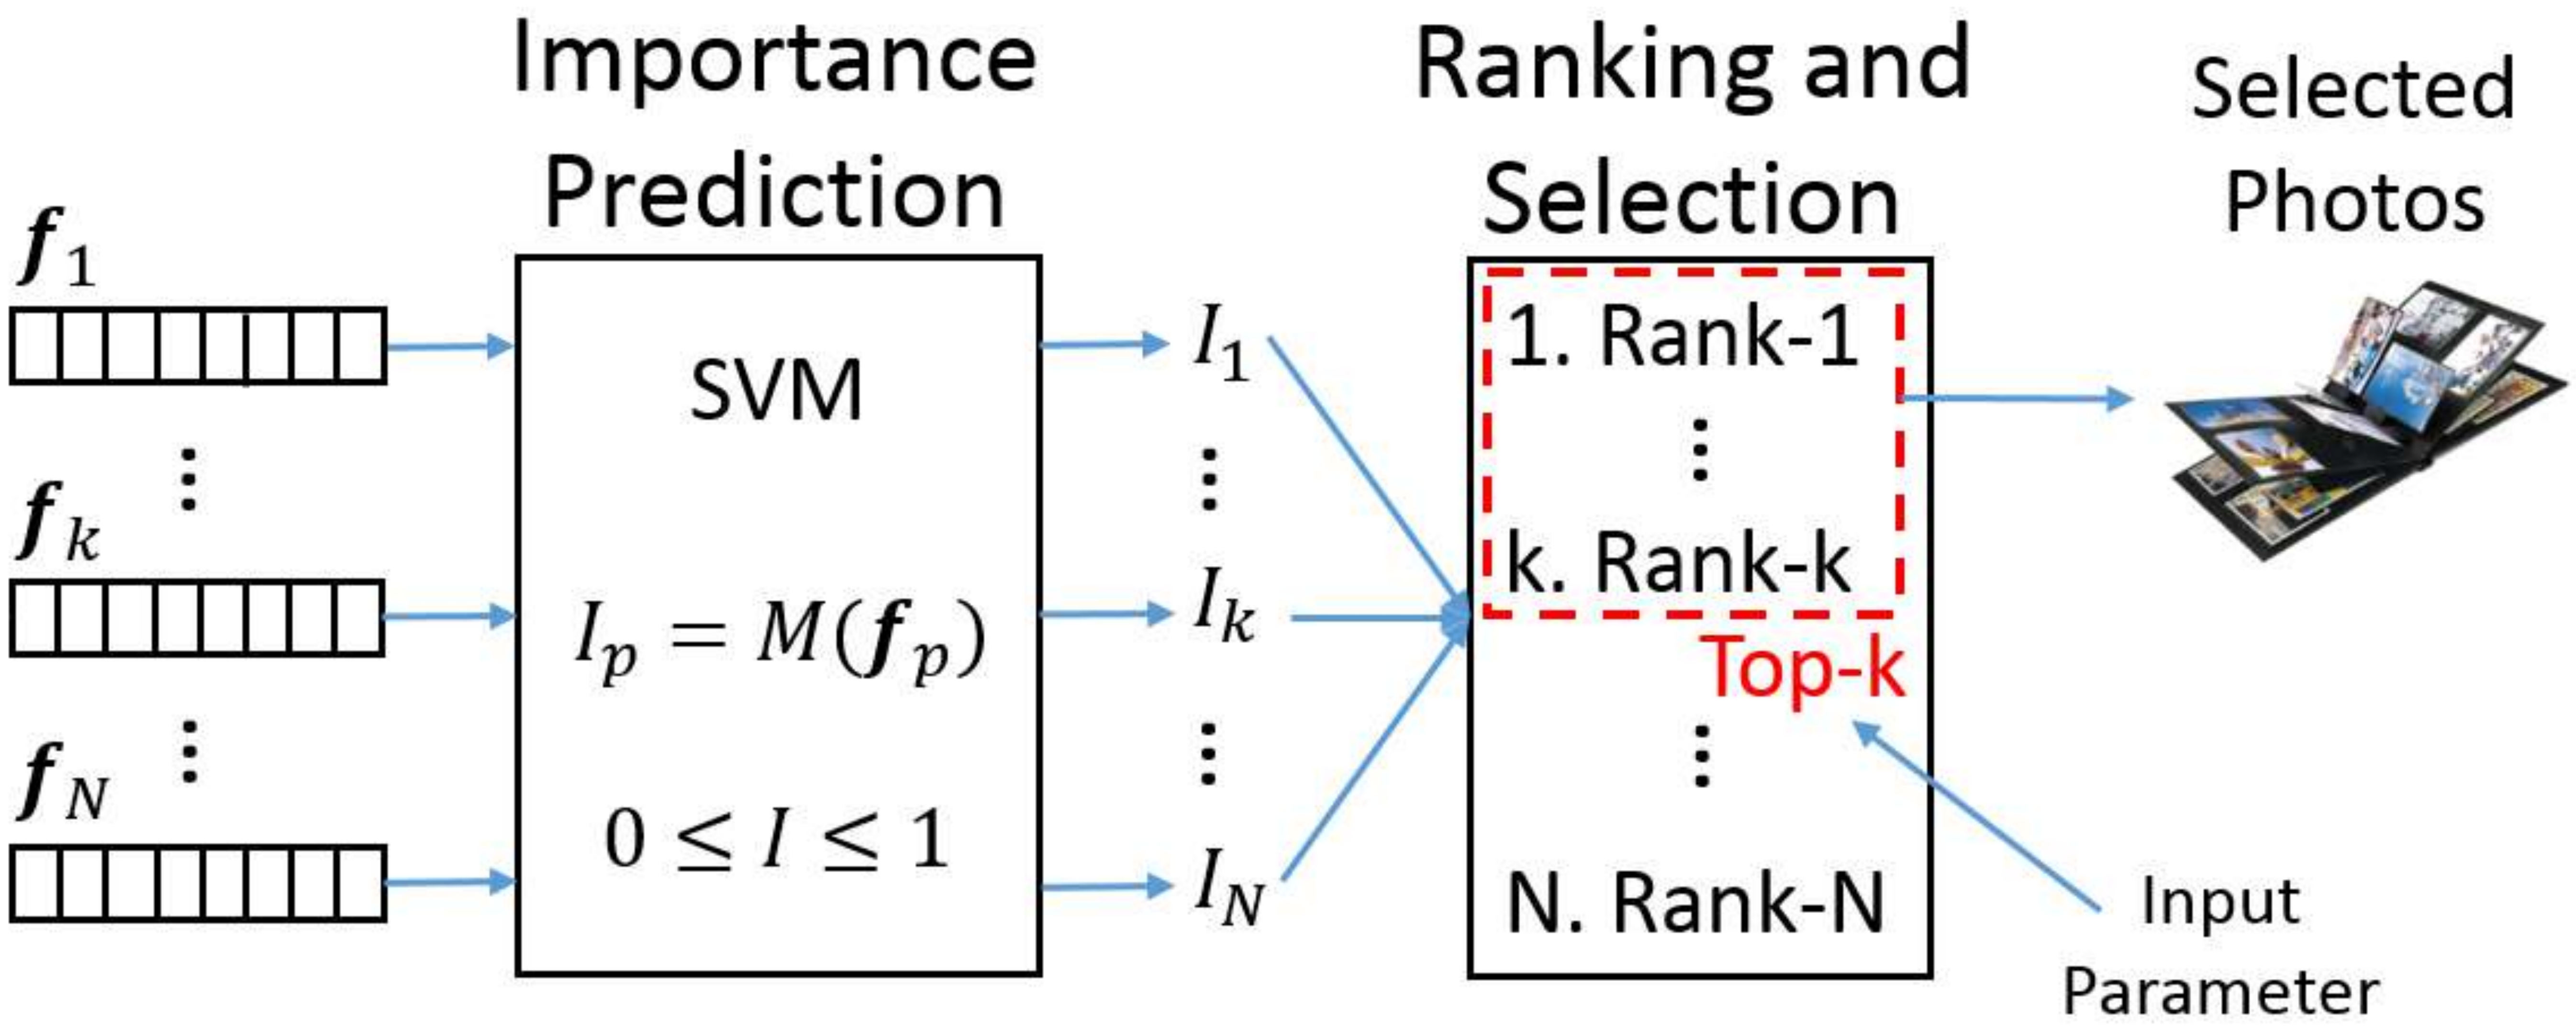

Experiments

- Ground Truth
  - Photo collections gathered during the user study
  - 91 collections, 42 users, 18,147 photos
  - 20% selected by the owner as most important for future revisiting

- Evaluation
  - Use the model to select the 20% photos from each collection
  - Count how many photos were also selected by the user (precision)

Results

✓ Statistically significant improvements over baselines

✓ Most important:

- Near-duplicates
- Faces
- Concepts

✗ Less important

- Image quality
- Clusters and sub-events

Other Directions

- What is the role of coverage and clustering in personal photo selection?
  - Attempts to incorporate coverage within the selection model
  - Coverage plays a secondary role in this task
- Inclusion of additional features in the model
  - Low-level visual info
  - Aesthetics
  - Deep Learning
  - Emotions
  - Face Clustering
- User Personalisation
  - Adapts to user preferences by exploiting user feedback

A. Ceroni. Methods for Managing, Validating, and Retrieving Event-related Information in Evolving Contexts. PhD Thesis, Leibniz Universität Hannover, 2018.

Thank you!

Contact info:

Andrea Ceroni
Data Scientist at SmartRecruiters GmbH
a.ceroni@smartrecruiters.com

Publications and Links:

- A. Ceroni et al.. To Keep or not to Keep: An Expectation-oriented Photo Selection Method for Personal Photo Collections. In ICMR 2015.
- A. Ceroni et al.. Investigating Human Behaviors in Selecting Personal Photos to Preserve Memories. In ICME Workshops 2015.
- A. Ceroni. (2018) Personal Photo Management and Preservation. In Personal Multimedia Preservation - Remembering or Forgetting Images and Videos. Springer, Cham.
- A. Ceroni. Methods for Managing, Validating, and Retrieving Event-related Information in Evolving Contexts. PhD Thesis, Leibniz Universität Hannover, 2018.
- ForgetIT Project: www.forgetit-project.eu

Backup Slides

Interface

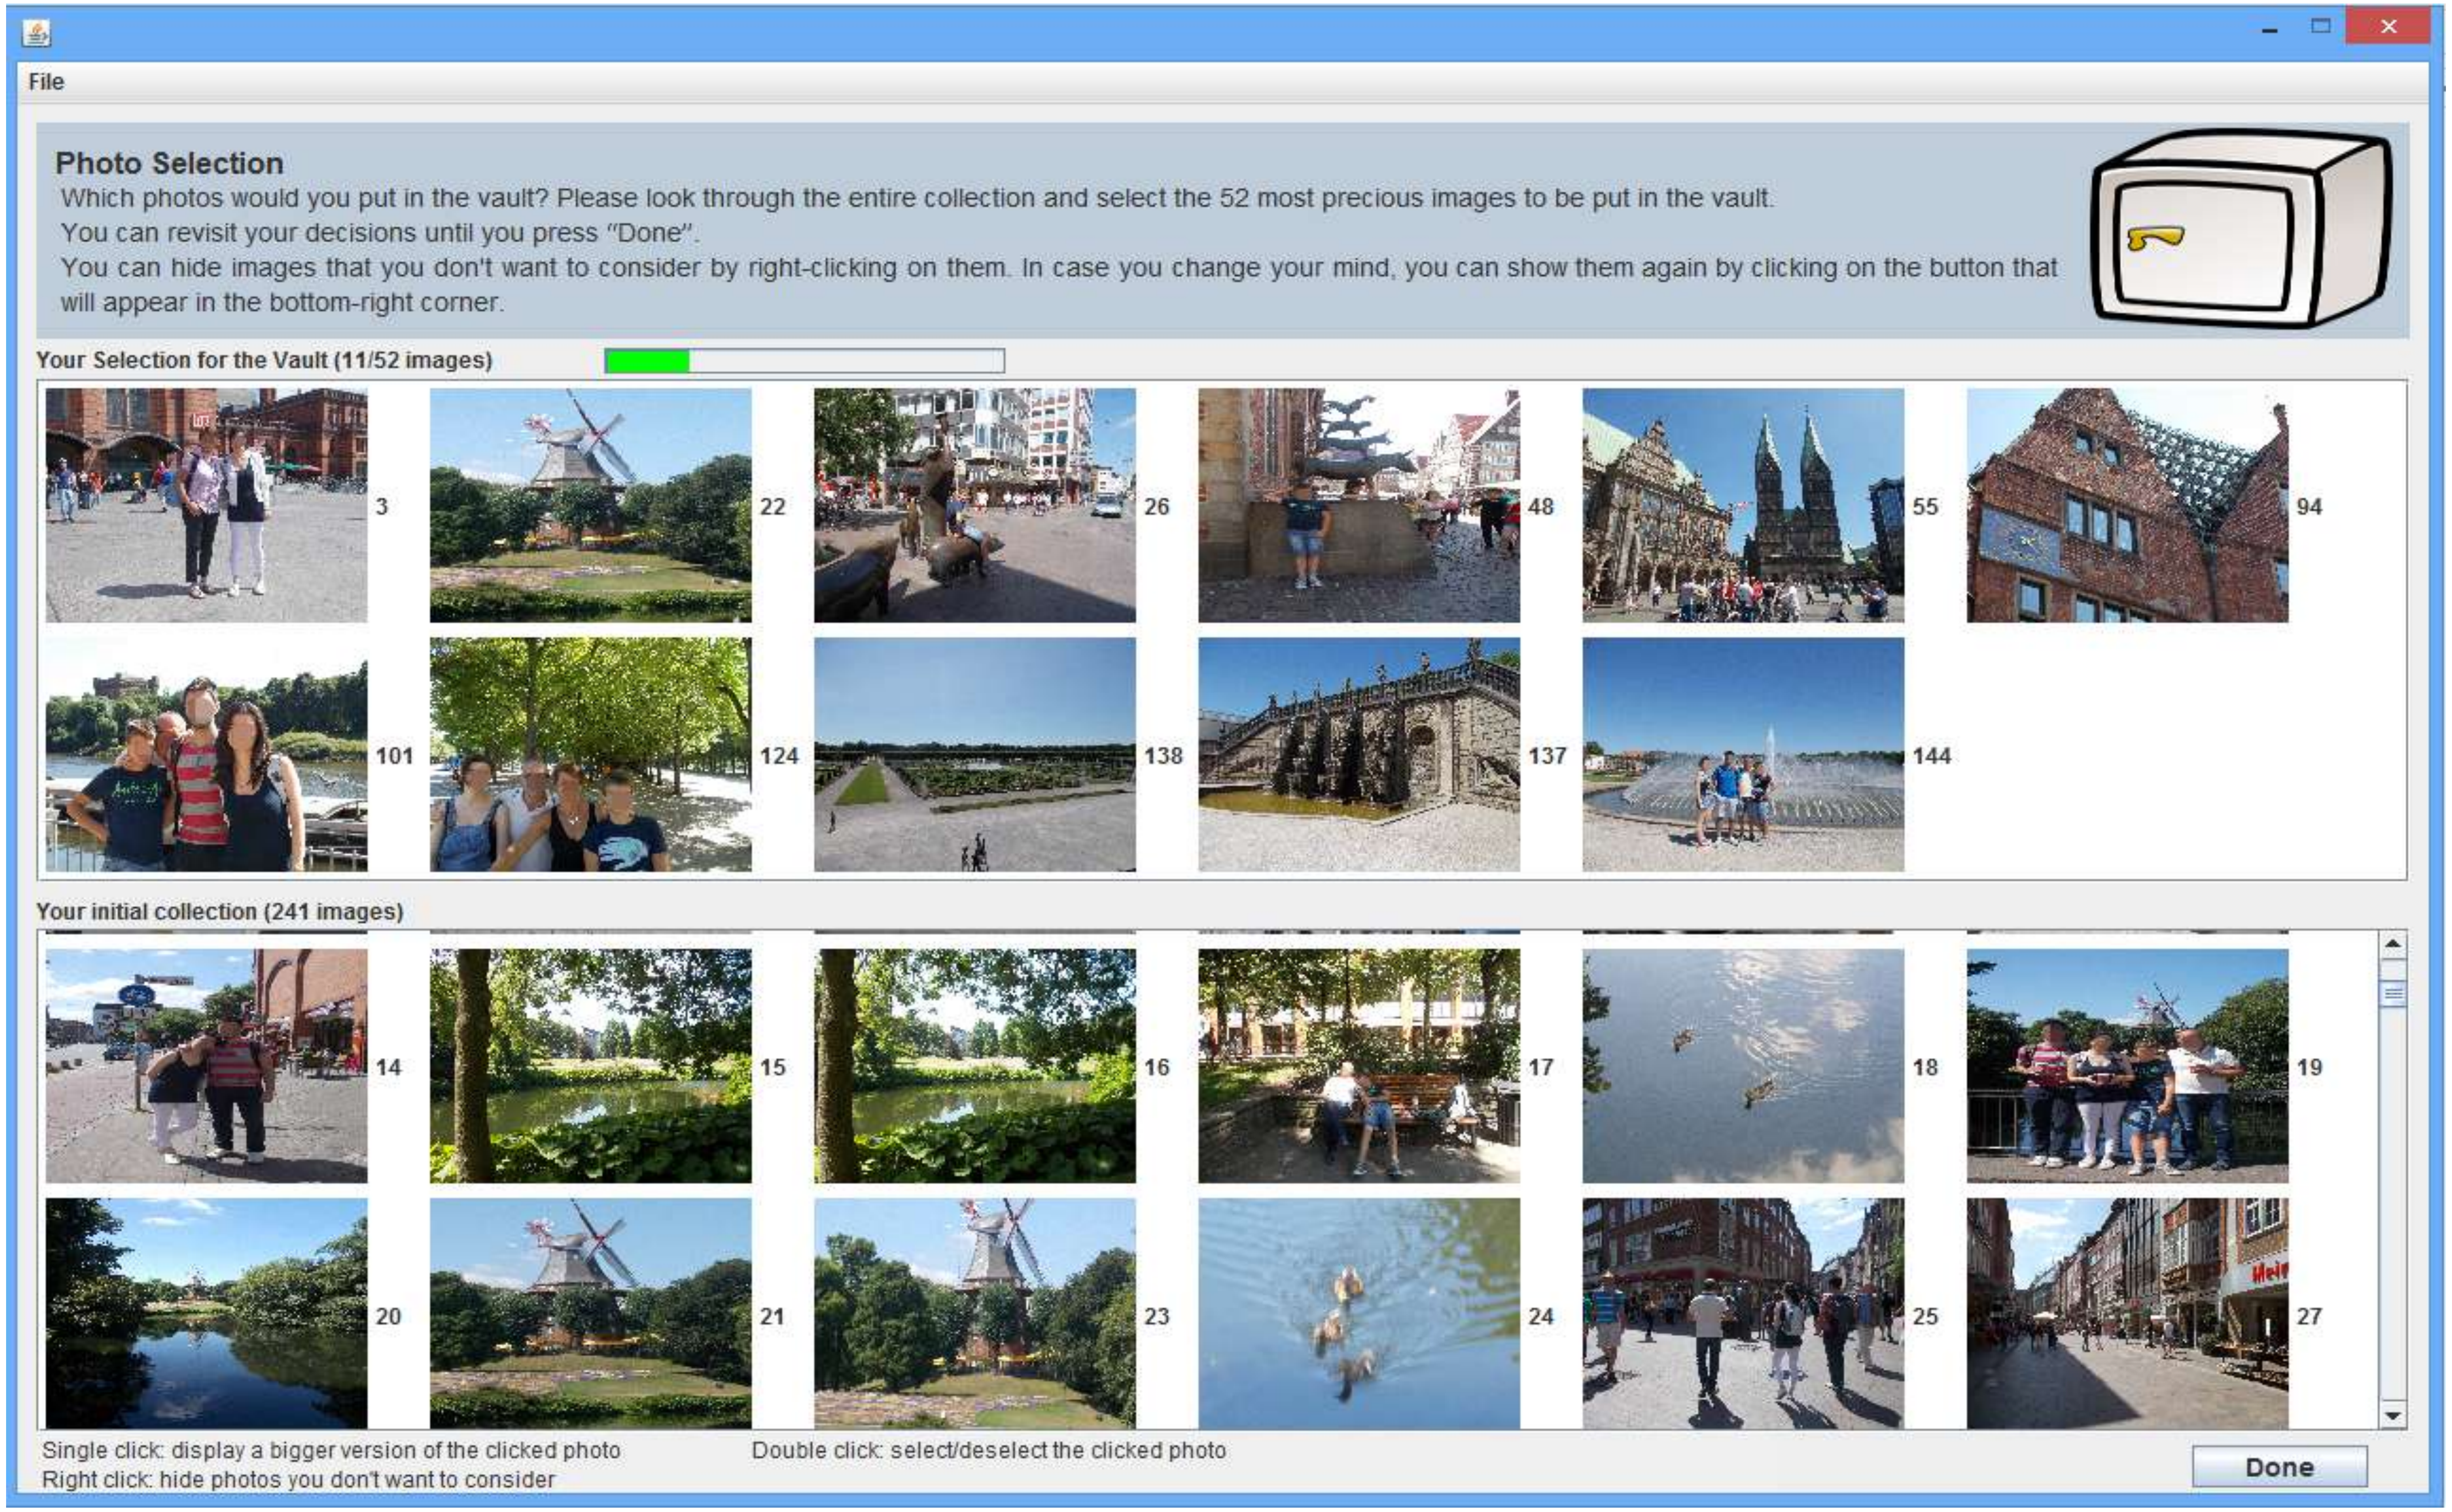
File
Photo Selection
Which photos would you put in the vault? Please look through the entire collection and select the 52 most precious images to be put in the vault.
You can revisit your decisions until you press "Done".
You can hide images that you don't want to consider by right-clicking on them. In case you change your mind, you can show them again by clicking on the button that will appear in the bottom-right corner.
Your Selection for the Vault (11/52 images)
3
22
26
48
55
94
101
124
138
137
144
Your initial collection (241 images)
14
15
16
17
18
19
20
21
23
24
25
27
Single click: display a bigger version of the clicked photo
Right click: hide photos you don't want to consider
Double click: select/deselect the clicked photo
Done

Automatic Photo Selection

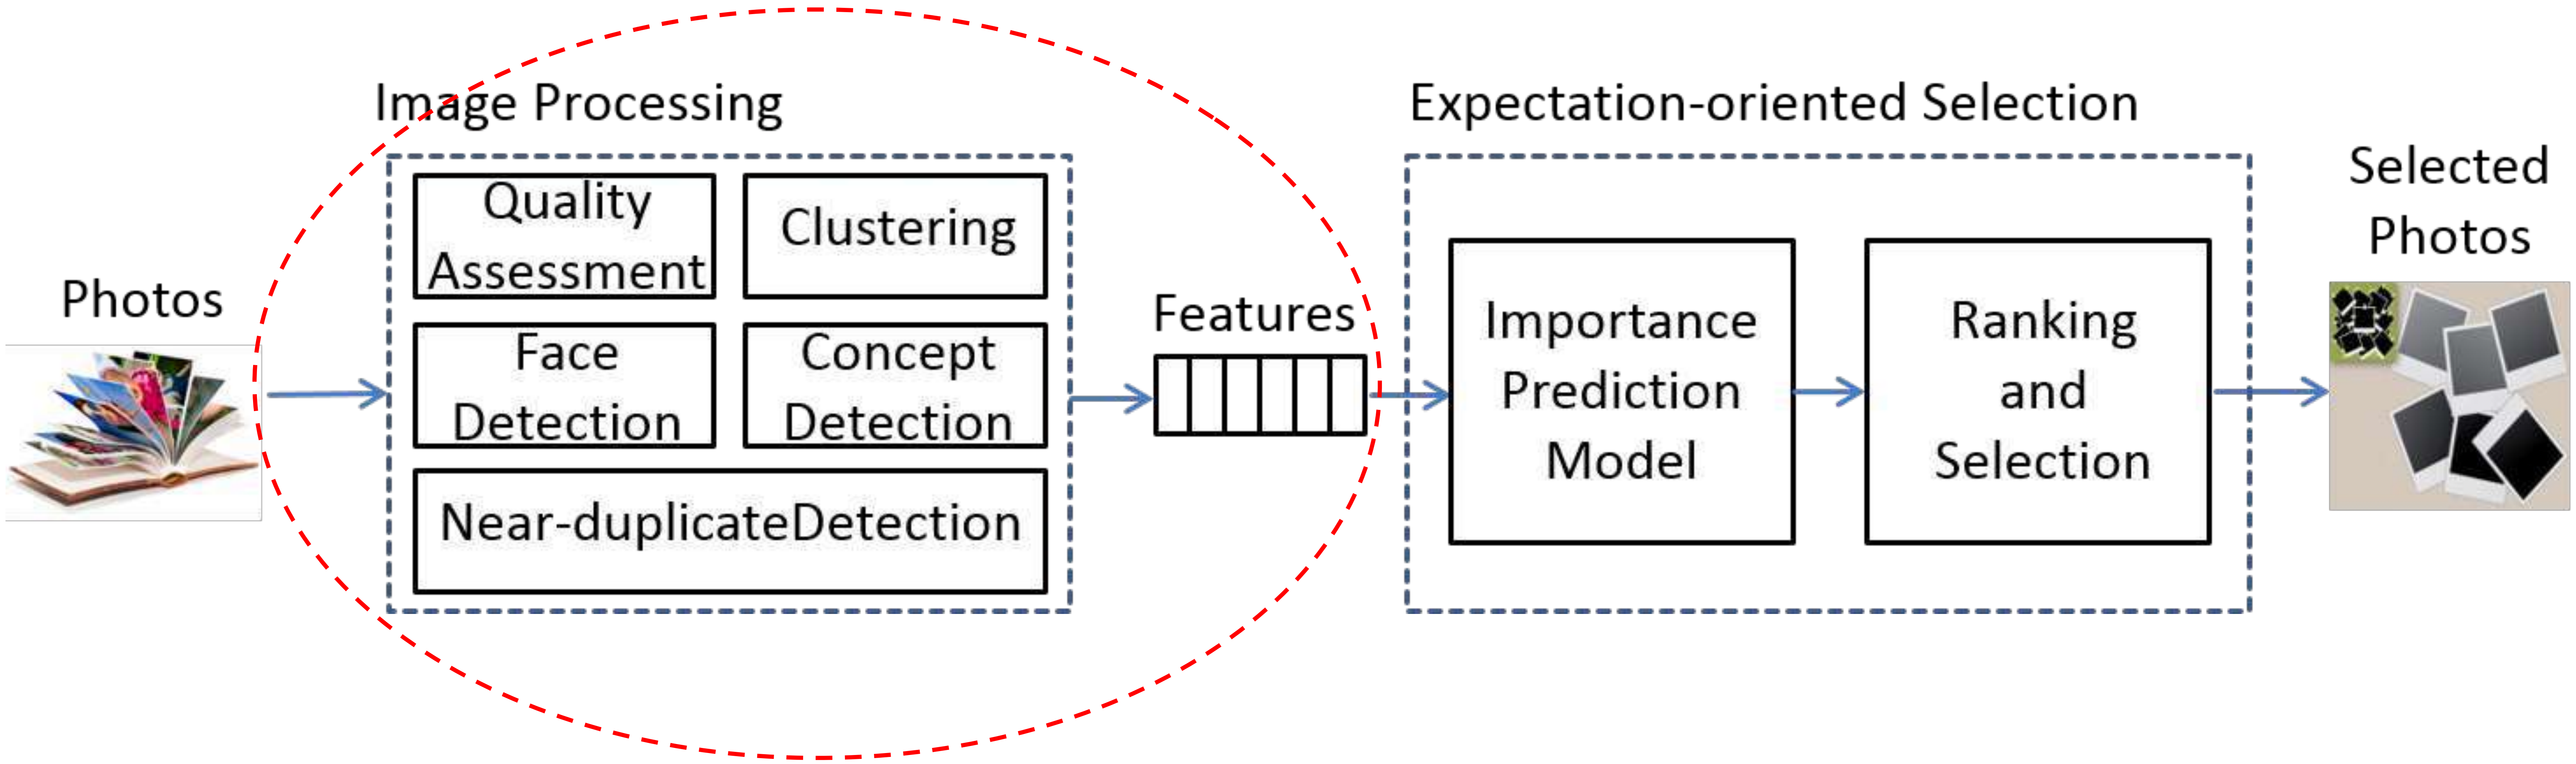

- ➢ User selections from personal collections used to train the model
- ➢ Relaxed notion of coverage (features from collections, clusters, near-duplicates)
- ➢ No manual annotations or external knowledge is required

Baselines

Temporal Clustering

- Cluster photos based on time [Cooper et al., 2005]
- Iterate the clusters (round robin)
- At each round, select the most important photo according to:

$$I(p) = \underbrace{\alpha \cdot \left\|\mathbf{q}_{\mathrm{p}}\right\|}_{\text{Quality}} + \underbrace{(1 - \alpha) \cdot \dim\left(\boldsymbol{F}_p\right)}_{\text{Faces}}, \quad \alpha = 0.3$$

Summary Optimization [Sinha et al., ICMR'11]

- Compute the optimal summary of size k according to:

$$S^* = \arg\max_{S \subset P_C} F(Qual(S),\ Div(S),\ Cov(S, P_C))$$

- Qual = sum of quality and *portrait, group, panorama* concepts values of each photo
- Div = diversity within the summary
- Cov = number of photos in the collection that are represented in the summary

Hybrid Selection

What is the role of coverage in personal photo selection?
Can we improve the selection by incorporating coverage within the model?

- Coverage-driven Selection
  - Cluster → Iterate → Select
  - Still a strict model of coverage

Importance Prediction

- Summary Optimization
  - Compute the optimal summary: $S^* = \arg\max_{S \subset P_C} F\left(Qual\left(S\right), Div\left(S\right), Cov\left(S, P_C\right)\right)$
  - More flexible

Results

Including importance prediction as quality measure in coverage-based methods improves their performances

A strict model of coverage via clustering gets smaller benefits

Expo is still better or comparable with the Hybrid Selection models

	P@5%	P@10%	P@15%	P@20%
Baselines				
Clustering	0.3741	0.3600	0.3436	0.3358
SummOpt	0.3858	0.3843	0.3687	0.3478
Coverage-driven Selection				
basic	0.4732▲	0.4113▲	0.3902△	0.3809△
filtered	0.5351▲	0.4617▲	0.4325▲	0.4170▲
filtered+greedy	0.6271▲	0.4835▲	0.4391▲	0.4262▲
SummOpt++	0.7115▲	0.5533▲	0.4937▲	0.4708▲
Expo	0.7124▲	0.5500▲	0.4895▲	0.4652▲

Statistically significant improvements marked as ▲ ($p < 0.01$) or Δ ($p < 0.05$).

Additional Features

Low-level visual info

Basic visual signals that might capture the attention and interest of the observer: HSV statistics, colors, textures, lines.

DCNN Features

Image representation given by a DCNN (GoogLeNet) pre-trained to predict the 1,000 categories of the ILSVRC.

Face Popularity

Face clustering applied to compute how frequently a face appears in a collection (cluster size).

Aesthetics

How an image is well posed, attractive and pleasant to an observer: rule of thirds, simplicity, contrast, balance.

Emotional Concepts

Concept detectors of SentiBank: nouns (concepts) and adjectives carrying sentiments are combined together to associate emotions to concepts.

Additional Features

Moderate yet statistically significant improvement

Concepts (**DCNN**) and concepts (**SentiBank**) improve **concepts** features

Face popularity only slightly improves **faces** features alone

Both **low level** and **aesthetics** features are better than **quality** features

	P@5%	**P@10%**	**P@15%**	**P@20%**
Expo				
quality	0.3431	0.3261	0.3204	0.3168
faces	0.4506	0.3968	0.3836	0.3747
concepts	0.5464	0.4599	0.4257	0.4117
all	**0.7124**	**0.5500**	**0.4895**	**0.4652**
Expo++				
low level	0.4399	0.3913	0.3729	0.3697
aesthetics	0.4406	0.3923	0.3732	0.3639
face popularity	0.4692	0.4101	0.3977	0.3945
concepts (DCNN)	0.5694	0.4945	0.4553	0.4436
concepts (SentiBank)	0.6124	0.5172	0.4674	0.4502
all	**0.7426**$^{\triangle}$	**0.6155**$^{\blacktriangle}$	**0.5330**$^{\blacktriangle}$	**0.5121**$^{\blacktriangle}$

User Personalization

Personalized photo selection model

- Adapts to user preferences by exploiting user feedback
- Based on retraining the model every time a new annotated collection is available

Promising adaptation capabilities

- Including new annotated collections of the same user can benefit future selections
- Exploiting annotated collections from other users can alleviate the cold-start problem

Evaluation on a large number of users and collections is required to make the results more evident and significant

Demo

- The presented application supports two use cases
  - Automatically selecting subsets of photos from personal photo collections for future revisiting, preservation, or sharing. The selection automatically done can be revised and modified by the user.
  - Gathering training data: users select the preferred photos manually and submit both the whole collection and the selection information.
- Novelty
  - The presented application selects photos according to the method presented in [1], which aims at predicting which photos the user perceives as most important and would select. It has been proved to be more effective in emulating user selections than methods based on coverage. It is based on a model trained via Machine Learning, considering information extracted through: concept detection, face detection, near-duplicate detection, quality and aesthetics assessment, event-based clustering.
  - Differently, available methods for summarization are usually centered around the concept of coverage, generating summaries that resemble the original collection. This is achieved either by clustering or explicitly modeling and optimizing coverage. We claim that selecting photos important to a user from a personal collection is a different task than generating comprehensive summaries, as the set of images important to a user might not totally resemble the original collection [1].

[1] Ceroni et al. To Keep or not to Keep: An Expectation–oriented Photo Selection Method for Personal Photo Collection. In ICMR '15